arXiv:1412.3978v4 [cs.GT] 24 Sep 2015

Delay Games with WMSO+U Winning Conditions*

Martin Zimmermann**

Reactive Systems Group, Saarland University, Germany
zimmermann@react.uni-saarland.de

Abstract. Delay games are two-player games of infinite duration in which one player may delay her moves to obtain a lookahead on her opponent's moves. We consider delay games with winning conditions expressed in weak monadic second order logic with the unbounding quantifier, which is able to express (un)boundedness properties.
We show that it is decidable whether the delaying player has a winning strategy using bounded lookahead and give a doubly-exponential upper bound on the necessary lookahead. In contrast, we show that bounded lookahead is not always sufficient to win such a game.

1 Introduction

Many of today's problems in computer science are no longer concerned with programs that transform data and then terminate, but with non-terminating reactive systems which have to interact with a possibly antagonistic environment for an unbounded amount of time. The framework of infinite two-player games is a powerful and flexible tool to verify and synthesize such systems. The seminal theorem of Büchi and Landweber [7] states that the winner of an infinite game on a finite arena with an ω-regular winning condition can be determined and a corresponding finite-state winning strategy can be constructed effectively.

Ever since, this result was extended along different dimensions, e.g., the number of players, the type of arena, the type of winning condition, the type of interaction between the players (alternation or concurrency), zero-sum or non-zero-sum, and complete or incomplete information. In this work, we consider two of these dimensions, namely more expressive winning conditions and the possibility for one player to delay her moves.

Delay Games. In a delay game, one of the players can postpone her moves for some time, thereby obtaining a lookahead on her opponent's moves. This allows her to win some games which she loses without lookahead, e.g., if her first move depends on the third move of her opponent. Nevertheless, there are winning conditions that cannot be won with any finite lookahead, e.g., if her first move

* A short version appears in the proceedings of CSR 2015 [31].
** Supported by the DFG projects "TriCS" (ZI 1516/1-1) and "AVACS" (SFB/TR 14).

depends on every move of her opponent. Delay arises naturally when transmission of data in networks or components equipped with buffers are modeled.

From a more theoretical point of view, uniformization of relations by continuous functions [27,28,29] can be expressed and analyzed using delay games. We consider games in which two players pick letters from alphabets Σ_I and Σ_O, respectively, thereby producing two infinite sequences $\alpha \in \Sigma_I^\omega$ and $\beta \in \Sigma_O^\omega$. Thus, a strategy for the second player induces a mapping $\tau\colon \Sigma_I^\omega \to \Sigma_O^\omega$. It is winning for the second player if $(\alpha, \tau(\alpha))$ is contained in the winning condition $L \subseteq \Sigma_I^\omega \times \Sigma_O^\omega$ for every α. If $\{(\alpha, \tau(\alpha)) \mid \alpha \in \Sigma_I^\omega\} \subseteq L$, then τ uniformizes L.

In the classical setting of infinite games, in which the players pick letters in alternation, the n-th letter of $\tau(\alpha)$ depends only on the first n letters of α, i.e., τ satisfies a very strong notion of continuity. A strategy with bounded lookahead, i.e., only finitely many moves are postponed, induces a Lipschitz-continuous function τ (in the Cantor topology on Σ^ω) and a strategy with arbitrary lookahead induces a continuous function (or equivalently, a uniformly continuous function, as Σ^ω is compact).

Hosch and Landweber proved that it is decidable whether a game with ω-regular winning condition can be won with bounded lookahead [19]. This result was improved by Holtmann, Kaiser, and Thomas who showed that if a player wins a game with arbitrary lookahead, then she wins already with doubly-exponential bounded lookahead, and gave a streamlined decidability proof yielding an algorithm with doubly-exponential running time [18]. Again, these results were improved by giving an exponential upper bound on the necessary lookahead and showing EXPTIME-completeness of the solution problem [20]. Going beyond ω-regular winning conditions by considering context-free conditions leads to undecidability and non-elementary lower bounds on the necessary lookahead, even for very weak fragments [15].

Thus, stated in terms of uniformization, Hosch and Landweber proved decidability of the uniformization problem for ω-regular relations by Lipschitz-continuous functions and Holtmann et al. proved the equivalence of the existence of a continuous uniformization function and the existence of a Lipschitz-continuous uniformization function for ω-regular relations. Furthermore, uniformization of context-free relations is undecidable, even with respect to Lipschitz-continuous functions.

In another line of work, Carayol and Löding considered the case of finite words [9], and Löding and Winter [23] considered the case of finite trees, which are both decidable. However, the nonexistence of MSO-definable choice functions on the infinite binary tree [8,17] implies that uniformization fails for such trees.

WMSO+U. In this work, we consider another class of conditions that go beyond the ω-regular ones. Recall that the ω-regular languages are exactly those that are definable in monadic second order logic (MSO) [6]. Recently, Bojańczyk has started a program investigating the logic MSO+U, MSO extended with the unbounding quantifier U. A formula $\mathrm{U}X\varphi(X)$ is satisfied, if there are arbitrarily large *finite* sets X such that $\varphi(X)$ holds. MSO+U is able to express all ω-regular

languages as well as non-regular languages like

$$L = \{a^{n_0}ba^{n_1}ba^{n_2}b\cdots \mid \limsup_i n_i = \infty\} \ .$$

Decidability of MSO+U turns out to be a delicate issue: there is no algorithm that decides MSO+U on infinite trees and has a correctness proof using the axioms of ZFC [3]. At the time of writing, an unconditional undecidability result for MSO+U on infinite words is presented [4].

Even before these undecidability results were shown, much attention was being paid to fragments of the logic obtained by restricting the power of the second-order quantifiers. In particular, considering weak[1] MSO with the unbounding quantifier (denoted by prepending a W) turned out to be promising: WMSO+U on infinite words [1] and on infinite trees [5] and WMSO+U with the path quantifier (WMSO+UP) on infinite trees [2] have equivalent automata models with decidable emptiness. Hence, these logics are decidable.

For WMSO+U on infinite words, these automata are called max-automata, deterministic automata with counters whose acceptance conditions are a boolean combination of conditions "counter c is bounded during the run". While processing the input, a counter may be incremented, reset to zero, or the maximum of two counters may be assigned to it (hence the name max-automata). In this work, we investigate delay games with winning conditions given by max-automata, so-called max-regular conditions.

Our Contribution. We prove the analogue of the Hosch-Landweber Theorem for max-regular winning conditions: it is decidable whether the delaying player has a winning strategy with bounded lookahead. Furthermore, we obtain a doubly-exponential upper bound on the necessary lookahead, if this is the case. Finally, we present a max-regular delay game such that the delaying player wins the game, but only with unbounded lookahead. Thus, unlike for ω-regular conditions, bounded lookahead is not sufficient for max-regular conditions. These are, to the best of our knowledge, the first results on delay games with quantitative winning conditions.

WMSO+U is able to express many quantitative winning conditions studied in the literature, e.g., winning conditions in parameterized temporal logics like Prompt-LTL [22], Parametric LTL [30], or Parametric LDL [13], finitary parity and Streett games [10], and parity and Streett games with costs [14]. Thus, for all these conditions we can decide whether Player O wins a delay game with bounded lookahead.

Our proof consists of a reduction to a delay-free game with a max-regular winning condition. Such games can be solved by expressing them as a satisfiability problem for WMSO+UP on infinite trees: the strategy of one player is an additional labeling of the tree and a path quantifier is able to range over all strategies of the opponent[2]. The reduction itself is an extension of the one used in the EXPTIME-algorithm for delay games with ω-regular winning conditions [20]

[1] Here, the second-order quantifiers are restricted to finite sets.
[2] See Example 1 in [2] for more details.

and is based on an equivalence relation that captures the behavior of the automaton recognizing the winning condition. However, unlike the relation used for ω-regular conditions, ours is only correct if applied to words of bounded lengths. Thus, we can deal with bounded lookahead, but not with arbitrary lookahead.

2 Definitions

The set of non-negative integers is denoted by $\mathbb{N}$. An alphabet Σ is a non-empty finite set of letters, and Σ^* (Σ^n, Σ^ω) denotes the set of finite words (words of length n, infinite words) over Σ. The empty word is denoted by ε, the length of a finite word w by $|w|$. For $w \in \Sigma^* \cup \Sigma^\omega$ we write $w(n)$ for the n-th letter of w.

Automata. Given a finite set C of counters storing non-negative integers,

$$\mathrm{Ops}(C) = \{c := c + 1, c := 0, c := \max(c_0, c_1) \mid c, c_0, c_1 \in C\}$$

is the set of counter operations over C. A counter valuation over C is a mapping $\nu\colon C \to \mathbb{N}$. By $\nu\pi$ we denote the counter valuation that is obtained by applying a finite sequence $\pi \in \mathrm{Ops}(C)^*$ of counter operations to ν, which is defined as implied by the operations' names.

A max-automaton $\mathcal{A} = (Q, C, \Sigma, q_I, \delta, \ell, \varphi)$ consists of a finite set Q of states, a finite set C of counters, an input alphabet Σ, an initial state q_I, a (deterministic and complete) transition function $\delta\colon Q \times \Sigma \to Q$, a transition labeling[3] $\ell\colon \delta \to \mathrm{Ops}(C)^*$ which labels each transition by a (possibly empty) sequence of counter operations, and an acceptance condition φ, which is a boolean formula over C.

A run of $\mathcal{A}$ on $\alpha \in \Sigma^\omega$ is an infinite sequence

$$\rho = (q_0, \alpha(0), q_1)\,(q_1, \alpha(1), q_2)\,(q_2, \alpha(2), q_3) \cdots \in \delta^\omega \tag{1}$$

with $q_0 = q_I$. Partial (finite) runs on finite words are defined analogously, i.e., $(q_0, \alpha(0), q_1) \cdots (q_{n-1}, \alpha(n-1), q_n)$ is the run of $\mathcal{A}$ on $\alpha(0) \cdots \alpha(n-1)$ starting in q_0. We say that this run ends in q_n. As δ is deterministic, $\mathcal{A}$ has a unique run on every finite or infinite word.

Let ρ be as in (1) and define $\pi_n = \ell(q_n, \alpha(n), q_{n+1})$, i.e., π_n is the label of the n-th transition of ρ. Given an initial counter valuation ν and a counter $c \in C$, we define the sequence

$$\rho_c = \nu(c)\,,\ \nu\pi_0(c)\,,\ \nu\pi_0\pi_1(c)\,,\ \nu\pi_0\pi_1\pi_2(c)\,, \ldots$$

of counter values of c reached on the run after applying *all* operations of a transition label. The run ρ of $\mathcal{A}$ on α is accepting, if the acceptance condition φ is satisfied by the variable valuation that maps a counter c to true if and only if $\limsup \rho_c$ is finite. Thus, φ can intuitively be understood as a boolean combination of conditions "$\limsup \rho_c < \infty$". Note that the limit superior of ρ_c is independent of the initial valuation used to define ρ_c, which is the reason it is

[3] Here, and later whenever convenient, we treat δ as relation $\delta \subseteq Q \times \Sigma \times Q$.

not part of the description of $\mathcal{A}$. We denote the language accepted by $\mathcal{A}$ by $L(\mathcal{A})$ and say that it is max-regular.

A parity condition (say min-parity) can be expressed in this framework using a counter for each color that is incremented every time this color is visited and employing the acceptance condition to check that the smallest color whose associated counter is unbounded, is even. Hence, the class of ω-regular languages is contained in the class of max-regular languages.

Given an automaton $\mathcal{A}$ over $\Sigma_I \times \Sigma_O$, we denote by $\pi_1(\mathcal{A})$ the automaton obtained by projecting each letter to its first component, which recognizes the projection of $L(\mathcal{A})$ to Σ_I.

Games with Delay. A delay function is a mapping $f\colon \mathbb{N} \to \mathbb{N} \setminus \{0\}$, which is said to be constant, if $f(i) = 1$ for every $i > 0$. Given a delay function f and an ω-language $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$, the game $\Gamma_f(L)$ is played by two players (Player I and Player O) in rounds $i = 0, 1, 2, \ldots$ as follows: in round i, Player I picks a word $u_i \in \Sigma_I^{f(i)}$, then Player O picks one letter $v_i \in \Sigma_O$. We refer to the sequence $(u_0, v_0), (u_1, v_1), (u_2, v_2), \ldots$ as a play of $\Gamma_f(L)$, which yields two infinite words $\alpha = u_0 u_1 u_2 \cdots$ and $\beta = v_0 v_1 v_2 \cdots$. Player O wins the play if and only if the outcome $\binom{\alpha(0)}{\beta(0)}\binom{\alpha(1)}{\beta(1)}\binom{\alpha(2)}{\beta(2)} \cdots$ is in L, otherwise Player I wins.

Given a delay function f, a strategy for Player I is a mapping $\tau_I\colon \Sigma_O^* \to \Sigma_I^*$ such that $|\tau_I(w)| = f(|w|)$, and a strategy for Player O is a mapping $\tau_O\colon \Sigma_I^* \to \Sigma_O$. Consider a play $(u_0, v_0), (u_1, v_1), (u_2, v_2), \ldots$ of $\Gamma_f(L)$. Such a play is consistent with τ_I, if $u_i = \tau_I(v_0 \cdots v_{i-1})$ for every i; it is consistent with τ_O, if $v_i = \tau_O(u_0 \cdots u_i)$ for every i. A strategy τ for Player p is winning for her, if every play that is consistent with τ is won by Player p. In this case, we say Player p wins $\Gamma_f(L)$. A delay game is determined, if one of the players has a winning strategy.

Theorem 1. *Delay games with max-regular winning conditions are determined.*

Proof. We model a delay game $\Gamma_f(L(\mathcal{A}))$ for a max-automaton $\mathcal{A}$ as a parity game[4] $\mathcal{G}$ with finitely many colors in a countable arena. As such games are determined [11,25], so is $\Gamma_f(L(\mathcal{A}))$.

A vertex of the parity game stores the round number $i \in \mathbb{N}$, an indicator $t \in \{I, O\}$ which denotes whose player's turn it is, a state q of $\mathcal{A}$, and the current lookahead $w \in \Sigma_I^*$. Furthermore, it stores the current counter valuation ν_{cur} and a counter valuation $\nu_{\max}$ keeping track of the maximal value that a counter has assumed thus far. Finally, for every counter c there is a boolean flag u_c that is set to true if the value $\nu_{\max}$ is updated.

The successors of a vertex of the form $(i, I, q, w, \nu_{\text{cur}}, \nu_{\max}, (u_c)_{c \in C})$, i.e., it is Player I's turn, have the form $(i, O, q, ww', \nu_{\text{cur}}, \nu_{\max}, (u_c)_{c \in C})$ for some $w' \in \Sigma_I^{f(i)}$, i.e., Player I makes his move in round i by picking some $w' \in \Sigma_I^{f(i)}$, which is appended to the current lookahead w. Dually, successors of a vertex of the form $(i, O, q, aw, \nu_{\text{cur}}, \nu_{\max}, (u_c)_{c \in C})$ for $a \in \Sigma_I$ and $w \in \Sigma_I^*$, i.e., it

[4] See, e.g., [16] for a detailed definition of parity games.

is Player O's turn, have the form $(i+1, I, \delta(q, \binom{a}{b}), w, \nu'_{\text{cur}}, \nu'_{\max}, (u'_c)_{c\in C})$ for some $b \in \Sigma_O$, where we have $\pi = \ell(q, \binom{a}{b}, \delta(q, \binom{a}{b}))$, $\nu'_{\text{cur}} = \nu_{\text{cur}}\pi$, $\nu'_{\max}(c) = \max(\nu'_{\text{cur}}(c), \nu_{\max}(c))$, and $u'_c = 1$ if and only if $\nu'_{\max}(c) > \nu_{\max}(c)$. Here, Player O makes her move in round i, which consists of picking a letter b. The state of $\mathcal{A}$, the variable valuations, and the flags are updated accordingly.

A play is winning for Player O, if the set of counters c whose flag u_c is set to 1 infinitely often, satisfies the winning condition φ. This is a Muller condition[5] defined on a finite set of colors, namely the powerset of the set of counters. Applying the LAR-reduction turns the Muller condition into a parity condition with finitely many colors while keeping the arena countable, which yields the parity game $\mathcal{G}$.

Let ν_0 map every counter to 0. Player p has a winning strategy for the delay game $\Gamma_f(L(\mathcal{A}))$ if and only if she has a winning strategy for the parity game $\mathcal{G}$ from the initial vertex $(0, I, q_I, \varepsilon, \nu_0, \nu_0, (0)_{c\in C})$. Thus, $\Gamma_f(L(\mathcal{A}))$ is determined, as $\mathcal{G}$ is determined. □

This result is also implied by a recent more general determinacy theorem for delay games with Borel winning conditions [21].

3 An Equivalence Relation for Max-Automata

Fix $\mathcal{A} = (Q, C, \Sigma, q_I, \delta, \ell, \varphi)$. We use notions introduced in [1] to define equivalence relations over sequences of counter operations and over words over Σ that capture the behavior of $\mathcal{A}$. To this end, we need to introduce some notation to deal with runs of $\mathcal{A}$. Given a state q and $w \in \Sigma^* \cup \Sigma^\omega$, let $\rho(q, w)$ be the run of $\mathcal{A}$ on w starting in q. If w is finite, then $\delta^*(q, w)$ denotes the state $\rho(q, w)$ ends in. The transition profile of $w \in \Sigma^*$ is the mapping $q \mapsto \delta^*(q, w)$.

First, we define inductively what it means for a sequence $\pi \in \text{Ops}(C)^*$ to transfer a counter c to a counter d. The empty sequence and the operation $c := c+1$ transfer every counter to itself. The operation $c := 0$ transfers every counter but c to itself and the operation $c := \max(c_0, c_1)$ transfers every counter but c to itself and transfers c_0 and c_1 to c. Furthermore, if π_0 transfers c to e and π_1 transfers e to d, then $\pi_0\pi_1$ transfers c to d. If π transfers c to d, then we have $\nu\pi(d) \geq \nu(c)$ for every counter valuation ν, i.e., the value of d after executing π is larger or equal to the value of c before executing π, independently of the initial counter values.

Furthermore, a sequence of counter operations π transfers c to d with an increment, if there is a counter e and a decomposition $\pi_0\,(e := e+1)\,\pi_1$ of π such that π_0 transfers c to e and π_1 transfers e to d. If π transfers c to d with an increment, then we have $\nu\pi(d) \geq \nu(c) + 1$ for every counter valuation ν.

Finally, we say that π is a c-trace of length m, if there is a decomposition $\pi = \pi_0 \cdots \pi_{m-1}$ and a sequence of counters $c_0, c_1, \ldots, c_m$ with $c_m = c$ such that each π_i transfers c_i to c_{i+1} with an increment. If π is a c-trace of length m, then we have $\nu\pi(c) \geq m$ for every counter valuation ν.

[5] Again, see [16] for a formal definition.

As only counter values reached after executing all counter operations of a transition label are considered in the semantics of max-automata, we treat $\Lambda = \{\ell(q, a, q') \mid (q, a, q') \in \delta\}$ as an alphabet. Every word $\lambda \in \Lambda^*$ can be flattened to a word in $\mathrm{Ops}(C)^*$, which is denoted by $\mathrm{flat}(\lambda)$. However, infixes, prefixes, or suffixes of λ are defined with respect to the alphabet Λ. We define $\ell(q, w) \in \Lambda^*$ to be the sequence of elements in Λ labeling the run $\rho(q, w)$.

Let ρ be a finite run of $\mathcal{A}$ and let $\pi \in \mathrm{Ops}(C)^*$. We say that ρ ends with π, if π is a suffix of $\mathrm{flat}(\ell(\rho))$. A finite or infinite run contains π, if it has a prefix that ends in π.

Lemma 1 ([1]). *Let ρ be a run of $\mathcal{A}$ and c a counter. Then, $\limsup \rho_c = \infty$ if and only if ρ contains arbitrarily long c-traces.*

We use the notions of transfer (with increment) to define the equivalence relations that capture $\mathcal{A}$'s behavior[6]. We say that $\lambda, \lambda' \in \Lambda^*$ are equivalent, denoted by $\lambda \equiv_{\mathrm{ops}} \lambda'$, if for all counters c and d:

1. the flattening of λ transfers c to d if and only if the flattening of λ' transfers c to d,
2. the flattening of λ transfers c to d with an increment if and only if the flattening of λ' transfers c to d with an increment, and
3. λ has a prefix whose flattening transfers c to d if and only if λ' has a prefix whose flattening transfers c to d.

Using this, we define two words $x, x' \in \Sigma^*$ to be equivalent, denoted by $x \equiv_{\mathcal{A}} x'$, if they have the same transition profile and if $\ell(q, x) \equiv_{\mathrm{ops}} \ell(q, x')$ for all states q.

Remark 1. Let $\mathcal{A}$ be a max-automaton with n states and k counters.

1. The index of $\equiv_{\mathrm{ops}}$ is at most 2^{k^2}.
2. The index of $\equiv_{\mathcal{A}}$ is at most $2^{n(\log n + 6k^2)}$.

Next, we show that we can decompose an infinite word α into $x_0 x_1 x_2 \cdots$ and replace each x_i by an $\equiv_{\mathcal{A}}$-equivalent x_i' without changing membership in $L(\mathcal{A})$, provided the lengths of the x_i and the lengths of the x_i' are bounded.

Lemma 2. *Let $(x_i)_{i \in \mathbb{N}}$ and $(x_i')_{i \in \mathbb{N}}$ be two sequences of words over Σ^* with $\sup_i |x_i| < \infty$, $\sup_i |x_i'| < \infty$, and $x_i \equiv_{\mathcal{A}} x_i'$ for all i. Then, $x = x_0 x_1 x_2 \cdots \in L(\mathcal{A})$ if and only if $x' = x_0' x_1' x_2' \cdots \in L(\mathcal{A})$.*

Proof. Let ρ and ρ' be the run of $\mathcal{A}$ on x and x', respectively. We show that ρ contains arbitrarily long c-traces if and only if ρ' contains arbitrarily long c-traces. Due to Lemma 1, this suffices to show that the run of $\mathcal{A}$ on x is accepting

[6] In the conference version of this paper [31], the third requirement is missing, which invalidates the proof of Lemma 2. The correction presented here has only a small influence on the bounds presented in Remark 1 and in Corollary 2, but the statement of Lemma 2 and thus also the proof of our main theorem are not affected.

if and only if the run of $\mathcal{A}$ on x' is accepting. Furthermore, due to symmetry, it suffices to show one direction of the equivalence. Thus, assume ρ contains arbitrarily long c-traces and pick $m' \in \mathbb{N}$ arbitrarily. We show the existence of a c-trace of length m' contained in ρ'. To this end, we take a c-trace in ρ of length $m > m'$ for some sufficiently large m and show that the $\equiv_{\text{ops}}$-equivalent part of ρ' contains a c-trace of length m'.

By definition of $\equiv_{\mathcal{A}}$, processing $x_0 \cdots x_{i-1}$ and processing $x'_0 \cdots x'_{i-1}$ brings $\mathcal{A}$ to the same state, call it q_i. Furthermore, let $\pi_i = \ell(q_i, x_i)$ be the sequence of counter operations labeling the run of $\mathcal{A}$ on x_i starting in q_i, which ends in q_{i+1}. The sequences π'_i labeling the runs on the x'_i are defined analogously. By $x_i \equiv_{\mathcal{A}} x'_i$ we conclude that π_i and π'_i are $\equiv_{\text{ops}}$-equivalent as well. Furthermore, define $b = \sup_i |x_i|$, which is well-defined due to our assumption, and define $m = (m'+1) \cdot o \cdot b$, where $o = \max_{\pi \in \Lambda} |\pi|$ is the maximal length of a sequence of operations labeling a transition. Each π_i can contribute at most $|\pi_i|$ increments to a c-trace that subsumes π_i, which is bounded by $|\pi_i| \leq o \cdot b$.

Now, we pick i such that $\pi_0 \cdots \pi_i$ contains a c-trace of length m. We can assume w.l.o.g. that the trace starts at the beginning of π_s for some $s \leq i$ and ends in a prefix of π_i. Hence, there are counters $c_s, c_{s+1}, \ldots, c_i$ such that the flattening of π_j transfers c_j to c_{j+1} for every j in the range $s \leq j < i$ and that π_i has a prefix whose flattening transfers c_i to c. Furthermore, by the choice of m we know that at least m' of these transfers are actually transfers with increments, as every transfer contains at most $b \cdot o$ increments.

The equivalence of π_j and π'_j implies that π'_j realizes the same transfers (with increments) as π_j. Hence, $\pi'_0 \cdots \pi'_i$ contains a c-trace of length m' as well. □

Note that the lemma does not hold if we drop the boundedness requirements on the lengths of the x_i and the x'_i.

To conclude, we show that the equivalence classes of $\equiv_{\mathcal{A}}$ are regular and can be *tracked* on-the-fly by a finite automaton $\mathcal{T}$ in the following sense.

Lemma 3. *There is a deterministic finite automaton $\mathcal{T}$ with set of states $\Sigma^*/\equiv_{\mathcal{A}}$ such that the run of $\mathcal{T}$ on $w \in \Sigma^*$ ends in $[w]_{\equiv_{\mathcal{A}}}$.*

Proof. Define $\mathcal{T} = (\Sigma^*/\equiv_{\mathcal{A}}, \Sigma, [\varepsilon]_{\equiv_{\mathcal{A}}}, \delta_{\mathcal{T}}, \emptyset)$ where $\delta_{\mathcal{T}}([x]_{\equiv_{\mathcal{A}}}, a) = [xa]_{\equiv_{\mathcal{A}}}$, which is independent of the representative x and based on the fact that $\equiv_{\text{ops}}$ (and thus also $\equiv_{\mathcal{A}}$) is a congruence, i.e., $\pi_0 \equiv_{\text{ops}} \pi_1$ implies $\pi_0\pi \equiv_{\text{ops}} \pi_1\pi$ for every π. A straightforward induction over $|w|$ shows that $\mathcal{T}$ has the desired properties. □

Corollary 1. *Every $\equiv_{\mathcal{A}}$-equivalence class is regular.*

4 Reducing Delay Games to Delay-free Games

In this section, we prove our main theorem.

Theorem 2. *The following problem is decidable: given a max-automaton $\mathcal{A}$, does Player O win $\Gamma_f(L(\mathcal{A}))$ for some constant delay function f?*

To prove this result, we construct a delay-free game in a finite arena with a max-regular winning condition that is won by Player O if and only if she wins $\Gamma_f(L(\mathcal{A}))$ for some constant delay function f. The winner of such a game can be determined effectively.

Let $\mathcal{A} = (Q, C, \Sigma_I \times \Sigma_O, q_I, \delta, \ell, \varphi)$ and let $\mathcal{T} = ((\Sigma_I \times \Sigma_O)/\equiv_{\mathcal{A}}, \Sigma_I \times \Sigma_O, [\varepsilon]_{\equiv_{\mathcal{A}}}, \delta_{\mathcal{T}}, \emptyset)$ be defined as in Lemma 3. For the sake of readability, we denote the $\equiv_{\mathcal{A}}$-equivalence class of w by $[w]$ without a subscript. Furthermore, we denote equivalence classes using the letter S. We define the product $\mathcal{P} = (Q_{\mathcal{P}}, C, \Sigma_I \times \Sigma_O, q_I^{\mathcal{P}}, \delta_{\mathcal{P}}, \ell_{\mathcal{P}}, \varphi)$ of $\mathcal{A}$ and $\mathcal{T}$, which is a max-automaton, where

- $Q_{\mathcal{P}} = Q \times ((\Sigma_I \times \Sigma_O)/\equiv_{\mathcal{A}})$,
- $q_I^{\mathcal{P}} = (q_I, [\varepsilon]_{\equiv_{\mathcal{A}}})$,
- $\delta_{\mathcal{P}}((q, S), a) = (\delta(q, a), \delta_{\mathcal{T}}(S, a))$ for a states $q \in Q$, an equivalence class $S \in (\Sigma_I \times \Sigma_O)/\equiv_{\mathcal{A}}$, and a letter $a \in \Sigma_I \times \Sigma_O$, and
- $\ell_{\mathcal{P}}((q, S), a, (q', S')) = \ell(q, a, q')$.

Let $n = |Q_{\mathcal{P}}|$. We have $L(\mathcal{P}) = L(\mathcal{A})$, since acceptance only depends on the component $\mathcal{A}$ of $\mathcal{P}$. However, we are interested in partial runs of $\mathcal{P}$, as the component $\mathcal{T}$ keeps track of the equivalence class of the input processed by $\mathcal{P}$.

Remark 2. Let $w \in (\Sigma_I \times \Sigma_O)^*$ and let $(q_0, S_0)(q_1, S_1) \cdots (q_{|w|}, S_{|w|})$ be the run of $\mathcal{P}$ on w from some state (q_0, S_0) with $S_0 = [\varepsilon]$. Then, $q_0 q_1 \cdots q_{|w|}$ is the run of $\mathcal{A}$ on w starting in q_0 and $S_{|w|} = [w]$.

In the following, we will work with partial functions r from $Q_{\mathcal{P}}$ to $2^{Q_{\mathcal{P}}}$, where we denote the domain of r by $\mathrm{dom}(r)$. Intuitively, we use such a function to capture the information encoded in the lookahead provided by Player I. Assume Player I has picked $\alpha(0) \cdots \alpha(j)$ and Player O has picked $\beta(0) \cdots \beta(i)$ for some $i < j$, i.e., the lookahead is $\alpha(i+1) \cdots \alpha(j)$. Then, we can determine the state q that $\mathcal{P}$ reaches when processing $\binom{\alpha(0)}{\beta(0)} \cdots \binom{\alpha(i)}{\beta(i)}$, but the automaton cannot process $\alpha(i+1) \cdots \alpha(j)$, since Player O has not yet provided her moves $\beta(i+1) \cdots \beta(j)$. However, we can determine which states Player O can enforce by picking an appropriate completion. These will be contained in $r(q)$.

To formalize this, we use the function $\delta_{\mathrm{pow}} \colon 2^{Q_{\mathcal{P}}} \times \Sigma_I \to 2^{Q_{\mathcal{P}}}$ defined via $\delta_{\mathrm{pow}}(P, a) = \bigcup_{q \in P} \bigcup_{b \in \Sigma_O} \delta_{\mathcal{P}}\left(q, \binom{a}{b}\right)$, i.e., δ_{pow} is the transition function of the powerset automaton of the projection automaton $\pi_1(\mathcal{P})$. As usual, we extend δ_{pow} to $\delta^*_{\mathrm{pow}} \colon 2^{Q_{\mathcal{P}}} \times \Sigma_I^* \to 2^{Q_{\mathcal{P}}}$ via $\delta^*_{\mathrm{pow}}(P, \varepsilon) = P$ and $\delta^*_{\mathrm{pow}}(P, wa) = \delta_{\mathrm{pow}}(\delta^*_{\mathrm{pow}}(P, w), a)$.

Let $D \subseteq Q_{\mathcal{P}}$ be non-empty and let $w \in \Sigma_I^*$. We define the function r_w^D with domain D as follows: for every $(q, S) \in D$, we have

$$r_w^D(q, S) = \delta^*_{\mathrm{pow}}(\{(q, [\varepsilon])\}, w) \ ,$$

i.e., we collect all states (q', S') reachable from $(q, [\varepsilon])$ (note that the second component is the equivalence class of the empty word, not the class S from the argument) via a run of $\pi_1(\mathcal{P})$ on w. Thus, if $(q', S') \in r_w^D(q, S)$, then there is a word w' whose projection is w and with $[w'] = S'$ such that the run of $\mathcal{A}$ on w'

leads from q to q'. Thus, if Player I has picked the lookahead w, then Player O could pick an answer such that the combined word leads $\mathcal{A}$ from q to q' and such that it is a representative of S'.

We call w a witness for a partial function $r\colon Q_{\mathcal{P}} \to 2^{Q_{\mathcal{P}}}$, if we have $r = r_w^{\mathrm{dom}(r)}$. Thus, we obtain a language $W_r \subseteq \Sigma_I^*$ of witnesses for each such function r. Now, we define $\mathfrak{R} = \{r \mid \mathrm{dom}(r) \neq \emptyset \text{ and } W_r \text{ is infinite}\}$.

Lemma 4. *Let $\mathfrak{R}$ be defined as above.*

1. *Let $r \in \mathfrak{R}$. Then, $r(q) \neq \emptyset$ for every $q \in \mathrm{dom}(r)$.*
2. *Let r be a partial function from $Q_{\mathcal{P}}$ to $2^{Q_{\mathcal{P}}}$. Then, W_r is recognized by a deterministic finite automaton with 2^{n^2} states.*
3. *Let $r \in \mathfrak{R}$. Then, W_r contains a word w with $k \leq |w| \leq k + 2^{n^2}$ for every k.*
4. *Let $r \neq r' \in \mathfrak{R}$ such that $\mathrm{dom}(r) = \mathrm{dom}(r')$. Then, $W_r \cap W_{r'} = \emptyset$.*
5. *Let $D \subseteq Q_{\mathcal{P}}$ be non-empty and let w be such that $|w| \geq 2^{n^2}$. Then, there exists some $r \in \mathfrak{R}$ with $\mathrm{dom}(r) = D$ and $w \in W_r$.*

Due to items 4.) and 5.), we can define for every non-empty $D \subseteq Q_{\mathcal{P}}$ a function r_D that maps words $w \in \Sigma_I^*$ with $|w| \geq 2^{n^2}$ to the unique function r with $\mathrm{dom}(r) = D$ and $w \in W_r$. This will be used later in the proof.

Now, we define an abstract game $\mathcal{G}(\mathcal{A})$ between Player I and Player O that is played in rounds $i = 0, 1, 2, \ldots$: in each round, Player I picks a function from $\mathfrak{R}$ and then Player O picks a state q of $\mathcal{P}$. In round 0, Player I has to pick r_0 subject to constraint (C1): $\mathrm{dom}(r_0) = \{q_I^{\mathcal{P}}\}$. Then, Player O has to pick a state $q_0 \in \mathrm{dom}(r_0)$ (which implies $q_0 = q_I^{\mathcal{P}}$). Now, consider round $i > 0$: Player I has picked functions $r_0, r_1, \ldots, r_{i-1}$ and Player O has picked states $q_0, q_1, \ldots, q_{i-1}$. Now, Player I has to pick a function r_i subject to constraint (C2): $\mathrm{dom}(r_i) = r_{i-1}(q_{i-1})$. Then, Player O has to pick a state $q_i \in \mathrm{dom}(r_i)$. Both players can always move: Player I can, as $r_{i-1}(q_{i-1})$ is always non-empty (Lemma 4.1) and thus the domain of some $r \in \mathfrak{R}$ (Lemma 4.5) and Player O can, as the domain of every $r \in \mathfrak{R}$ is non-empty by construction.

The resulting play is the sequence $r_0 q_0 r_1 q_1 r_2 q_2 \cdots$. Let $q_i = (q_i', S_i)$ for every i, i.e., S_i is an $\equiv_{\mathcal{A}}$-equivalence class. Let $x_i \in S_i$ for every i such that $\sup_i |x_i| < \infty$. Such a sequence can always be found as $\equiv_{\mathcal{A}}$ has finite index. Player O wins the play if the word $x_0 x_1 x_2 \cdots$ is accepted by $\mathcal{A}$. Due to Lemma 2, this definition is independent of the choice of the representatives x_i.

A strategy for Player I is a function τ_I' mapping the empty play prefix to a function r_0 subject to constraint (C1) and mapping a non-empty play prefix $r_0 q_0 \cdots r_{i-1} q_{i-1}$ ending in a state to a function r_i subject to constraint (C2). On the other hand, a strategy for Player O maps a play prefix $r_0 q_0 \cdots r_i$ ending in a function to a state $q_i \in \mathrm{dom}(r_i)$. A play $r_0 q_0 r_1 q_1 r_2 q_2 \cdots$ is consistent with τ_I', if $r_i = \tau_I'(r_0 q_0 \cdots r_{i-1} q_{i-1})$ for every $i \geq 0$. Dually, the play is consistent with τ_O', if $q_i = \tau_O'(r_0 q_0 \cdots r_i)$ for every $i \geq 0$. A strategy is winning for Player p, if every play that is consistent with this strategy is winning for her. As usual, we say that Player p wins $\mathcal{G}(\mathcal{A})$, if she has a winning strategy.

Lemma 5. *Player O wins $\Gamma_f(L(\mathcal{A}))$ for some constant delay function f if and only if Player O wins $\mathcal{G}(\mathcal{A})$.*

Proof. First, assume Player O has a winning strategy τ_O for $\Gamma_f(L(\mathcal{A}))$ for some constant delay function f. We construct a winning strategy τ'_O for Player O in $\mathcal{G}(\mathcal{A})$ via simulating a play of $\mathcal{G}(\mathcal{A})$ by a play of $\Gamma_f(L(\mathcal{A}))$.

Let r_0 be the first move of Player I in $\mathcal{G}(\mathcal{A})$, which has to be responded to by Player O by picking $q_I^{\mathcal{P}} = \tau'_O(r_0)$, and let r_1 be Player I's response to that move. Let $w_0 \in W_{r_0}$ and $w_1 \in W_{r_1}$ be witnesses for the functions picked by Player I. Due to Lemma 4.3, we can choose w_0 and $|w_1|$ with $f(0) \leq |w_0|, |w_1| \leq f(0)+2^{n^2}$.

We simulate the play prefix $r_0 q_0 r_1$ in $\Gamma_f(L(\mathcal{A}))$, where $q_0 = q_I^{\mathcal{P}}$: Player I picks $w_0 w_1 = \alpha(0)\cdots\alpha(\ell_1 - 1)$ in his first moves and let $\beta(0)\cdots\beta(\ell_1 - f(0))$ be the response of Player O according to τ_O. We obtain $|\beta(0)\cdots\beta(\ell_1 - f(0))| \geq |w_0|$, as $|w_1| \geq f(0)$.

Thus, we are in the following situation for $i = 1$: in $\mathcal{G}(\mathcal{A})$, we have a play prefix $r_0 q_0 \cdots r_{i-1} q_{i-1} r_i$ and in $\Gamma_f(L(\mathcal{A}))$, Player I has picked $w_0 w_1 \cdots w_i = \alpha(0)\cdots\alpha(\ell_i - 1)$ and Player O has picked $\beta(0)\cdots\beta(\ell_i - f(0))$ according to τ_O, where $|\beta(0)\cdots\beta(\ell_i - f(0))| \geq |w_0 \cdots w_{i-1}|$. Furthermore, w_j is a witness for r_j for every $j \leq i$.

In this situation, let q_i be the state of $\mathcal{P}$ that is reached when processing w_{i-1} and the corresponding moves of Player O, i.e.,

$$\binom{\alpha(|w_0 \cdots w_{i-2}|)}{\beta(|w_0 \cdots w_{i-2}|)} \cdots \binom{\alpha(|w_0 \cdots w_{i-1}| - 1)}{\beta(|w_0 \cdots w_{i-1}| - 1)} ,$$

starting in state $(q'_{i-1}, [\varepsilon])$, where $q_{i-1} = (q'_{i-1}, S_{i-1})$.

By definition of r_{i-1}, we have $q_i \in r_{i-1}(q_{i-1})$, i.e., q_i is a legal move for Player O in $\mathcal{G}(\mathcal{A})$ to extend the play prefix $r_0 q_0 \cdots r_{i-1} q_{i-1} r_i$. Thus, we define $\tau'_O(r_0 q_0 \cdots r_{i-1} q_{i-1} r_i) = q_i$. Now, let r_{i+1} be the next move of Player I in $\mathcal{G}(\mathcal{A})$ and let $w_{i+1} \in W_{r_{i+1}}$ be a witness with $f(0) \leq |w_{i+1}| \leq f(0) + 2^{n^2}$. Going back to $\Gamma_f(L(\mathcal{A}))$, let Player I pick $w_{i+1} = \alpha(\ell_i)\cdots\alpha(\ell_{i+1} - 1)$ as his next moves and let $\beta(\ell_i - f(0) + 1)\cdots\beta(\ell_{i+1} - f(0))$ be the response of Player O according to τ_O. Then, we are in the situation as described in the previous paragraph, which concludes the definition of τ'_O.

It remains to show that τ'_O is winning for Player O in $\mathcal{G}(\mathcal{A})$. Consider a play $r_0 q_0 r_1 q_1 r_2 q_2 \cdots$ that is consistent with τ'_O and let $w = \binom{\alpha(0)}{\beta(0)}\binom{\alpha(1)}{\beta(1)}\binom{\alpha(2)}{\beta(2)}\cdots$ be the corresponding outcome constructed as in the simulation described above. Let $q_i = (q'_i, S_i)$, i.e., q'_i is a state of our original automaton $\mathcal{A}$. A straightforward inductive application of Remark 2 shows that q'_i is the state that $\mathcal{A}$ reaches after processing w_i and the corresponding moves of Player O, i.e.,

$$x_i = \binom{\alpha(|w_0 \cdots w_{i-1}|)}{\beta(|w_0 \cdots w_{i-1}|)} \cdots \binom{\alpha(|w_0 \cdots w_i| - 1)}{\beta(|w_0 \cdots w_i| - 1)} ,$$

starting in q'_{i-1}, and that $S_i = [x_i]$. Note that the length of the x_i is bounded, i.e., we have $\sup_i |x_i| \leq f(0) + 2^{n^2}$.

As w is consistent with a winning strategy for Player O, the run of $\mathcal{A}$ on $w = x_0 x_1 x_2 \cdots$ is accepting. Thus, we conclude that the play $r_0 q_0 r_1 q_1 r_2 q_2 \cdots$ is winning for Player O, as the x_i are a bounded sequence of representatives. Hence, τ'_O is indeed a winning strategy for Player O in $\mathcal{G}(\mathcal{A})$.

Now, we consider the other implication: assume Player O has a winning strategy τ'_O for $\mathcal{G}(\mathcal{A})$ and fix $d = 2^{n^2}$. We construct a winning strategy τ_O for her in $\Gamma_f(L(\mathcal{A}))$ for the constant delay function f with $f(0) = 2d$. In the following, both players pick their moves in blocks of length d. We denote Player I's blocks by a_i and Player O's blocks by b_i, i.e., in the following, every a_i is in Σ_I^d and every b_i is in Σ_O^d. This time, we simulate a play of $\Gamma_f(L(\mathcal{A}))$ by a play in $\mathcal{G}(\mathcal{A})$.

Let a_0a_1 be the first move of Player I in $\Gamma_f(L(\mathcal{A}))$, let $q_0 = q_I^{\mathcal{P}}$, and define the functions $r_0 = r_{\{q_0\}}(a_0)$ and $r_1 = r_{r_0(q_0)}(a_1)$. Then, $r_0q_0r_1$ is a legal play prefix of $\mathcal{G}(\mathcal{A})$ that is consistent with the winning strategy τ'_O for Player O.

Thus, we are in the following situation for $i = 1$: in $\mathcal{G}(\mathcal{A})$, we have constructed a play prefix $r_0q_0 \cdots r_{i-1}q_{i-1}r_i$ that is consistent with τ'_O; in $\Gamma_f(L(\mathcal{A}))$, Player I has picked $a_0 \cdots a_i$ such that a_j is a witness for r_j for every j in the range $0 \leq j \leq i$. Player O has picked $b_0 \cdots b_{i-2}$, which is the empty word for $i = 1$.

In this situation, let $q_i = \tau'_O(r_0q_0 \cdots r_{i-1}q_{i-1}r_i)$. By definition, we have $q_i \in \mathrm{dom}(r_i) = r_{i-1}(q_{i-1})$. Furthermore, as a_{i-1} is a witness for r_{i-1}, there exists b_{i-1} such that $\mathcal{P}$ reaches the state q_i when processing $\binom{a_{i-1}}{b_{i-1}}$ starting in state $(q'_{i-1}, [\varepsilon])$, where $q_{i-1} = (q'_{i-1}, S_{i-1})$.

Player O's strategy for $\Gamma_f(L(\mathcal{A}))$ is to play b_{i-1} in the next d rounds, which is answered by Player I by picking some a_{i+1} during these rounds. This induces the function $r_{i+1} = r_{r_i(q_i)}(a_{i+1})$. Now, we are in the same situation as described in the previous paragraph. This finishes the description of the strategy τ_O for Player O in $\Gamma_f(L(\mathcal{A}))$.

It remains to show that τ_O is winning for Player O. Let $w = \binom{a_0}{b_0}\binom{a_1}{b_1}\binom{a_2}{b_2}\cdots$ be the outcome of a play in $\Gamma_f(L(\mathcal{A}))$ that is consistent with τ_O. Furthermore, let $r_0q_0r_1q_1r_2q_2\cdots$ be the corresponding play in $\mathcal{G}(\mathcal{A})$ constructed in the simulation as described above, which is consistent with τ'_O. Let $q_i = (q'_i, S_i)$. A straightforward inductive application of Remark 2 shows that q'_i is the state reached by $\mathcal{A}$ after processing $x_i = \binom{a_i}{b_i}$ starting in q'_{i-1} and $S_i = [x_i]$. Furthermore, we have $\sup_i |x_i| = d$.

As $r_0q_0r_1q_1r_2q_2\cdots$ is consistent with a winning strategy for Player O and therefore winning for Player O, we conclude that $x_0x_1x_2\cdots$ is accepted by $\mathcal{A}$. Hence, $\mathcal{A}$ accepts the outcome w, which is equal to $x_0x_1x_2\cdots$, i.e., the play in $\Gamma_f(L(\mathcal{A}))$ is winning for Player O. Thus, τ_O is a winning strategy for Player O in $\Gamma_f(L(\mathcal{A}))$. □

Now, we can prove our main theorem of this section, Theorem 2.

Proof. Due to Lemma 5, we just have to show that we can construct and solve an explicit version of $\mathcal{G}(\mathcal{A})$. To this end, we encode $\mathcal{G}(\mathcal{A})$ as a graph-based game with arena (V, V_I, V_O, E) where

- the set of vertices is $V = V_I \cup V_O$ with
- the vertices $V_I = \{v_I\} \cup \mathfrak{R} \times Q_{\mathcal{P}}$ of Player I, where v_I is a fresh initial vertex,
- the vertices $V_O = \mathfrak{R}$ of Player O, and
- E is the union of the following sets of edges:
  - $\{(v_I, r) \mid \mathrm{dom}(r) = \{q_I^{\mathcal{P}}\}\}$: the initial moves of Player I.
  - $\{((r, q), r') \mid \mathrm{dom}(r') = r(q)\}$: (regular) moves of Player I.

- $\{(r,(r,q)) \mid q \in \mathrm{dom}(r)\}$: moves of Player O.

A play is an infinite path starting in v_I. To determine the winner of a play, we fix an arbitrary function rep: $(\Sigma_I \times \Sigma_O)^*/\equiv_{\mathcal{A}} \to (\Sigma_I \times \Sigma_O)^*$ that maps each equivalence class to some representative, i.e., $\mathrm{rep}(S) \in S$ for every $S \in (\Sigma_I \times \Sigma_O)^*/\equiv_{\mathcal{A}}$. Consider an infinite play

$$v_I, r_0, (r_0, q_0), r_1, (r_1, q_1), r_2, (r_2, q_2), \dots \ ,$$

with $q_i = (q_i', S_i)$ for every i. This play is winning for Player O, if the infinite word $\mathrm{rep}(S_0)\mathrm{rep}(S_1)\mathrm{rep}(S_2)\cdots$ is accepted by $\mathcal{A}$ (note that $\sup_i |\mathrm{rep}(S_i)|$ is bounded, as there are only finitely many equivalence classes). The set $\mathrm{Win} \subseteq V^\omega$ of winning plays for Player O is a max-regular language[7], as it can be recognized by an automaton that simulates the run of $\mathcal{A}$ on $\mathrm{rep}(S)$ when processing a vertex of the form $(r,(q,S))$ and ignores all other vertices. Games in finite arenas with max-regular winning condition are decidable via an encoding as a satisfiability problem for WMSO+UP [2].

Player O wins $\mathcal{G}(\mathcal{A})$ (and thus $\Gamma_f(L(\mathcal{A}))$ for some constant f) if and only if she has a winning strategy from v_I in the game $((V, V_I, V_O, E), \mathrm{Win})$, which concludes the proof. □

We obtain a doubly-exponential upper bound on the constant delay necessary for Player O to win a delay game with a max-regular winning condition by applying both directions of the equivalence between $\Gamma_f(\mathcal{A})$ and $\mathcal{G}(\mathcal{A})$.

Corollary 2. *Let $\mathcal{A}$ be a max-automaton with n states and k counters. The following are equivalent:*

1. *Player O wins $\Gamma_f(L(\mathcal{A}))$ for some constant delay function f.*
2. *Player O wins $\Gamma_f(L(\mathcal{A}))$ for some constant delay function f with $f(0) \leq 2^{2^{2n(\log n + 6k^2)}+1}$.*

5 Constant Delay Does Not Suffice

In this section, we show that constant delay does not suffice to win every delay game that Player O can with with arbitrary delay, i.e., the analogue of the Holtmann-Kaiser-Thomas theorem for delay games with max-regular winning conditions does not hold. In terms of uniformization, the following theorem shows that there are max-regular languages that are uniformizable by (uniformly) continuous functions, but not by Lipschitz-continuous functions.

Theorem 3. *There is a max-regular language L such that Player O wins $\Gamma_f(L)$ for some f, but not for any constant f.*

[7] This implies that $\mathcal{G}(\mathcal{A})$ is determined, as max-regular conditions are Borel [1,24].

Proof. Let $\Sigma_I = \{0, 1, \#\}$ and $\Sigma_O = \{0, 1, *\}$. An input block is a word $\#w$ with $w \in \{0,1\}^+$, and its length is defined to be $|w|$. An output block is a word

$$\binom{\#}{b}\binom{\alpha(1)}{*}\binom{\alpha(2)}{*}\cdots\binom{\alpha(n)}{*}\binom{b}{b} \in (\Sigma_I \times \Sigma_O)^+$$

for $b \in \{0,1\}$ and $\alpha(j) \in \{0,1\}$ for all j in the range $1 \leq j \leq n$. The length of the block is $n+1$. Note that the first and last letter in an output block are the only ones whose second component is not a $*$, and that these letters have to be equal to the first component of the last letter of the block. Every input block of length n can be extended to an output block of length n and and the projection to the first component of every output block is an input block.

A word $\binom{\alpha(0)}{\beta(0)}\binom{\alpha(1)}{\beta(1)}\binom{\alpha(2)}{\beta(2)}\cdots$ over $\Sigma_I \times \Sigma_O$ is in L if and only if it satisfies the following property: if $\alpha(0)\alpha(1)\alpha(2)\cdots$ contains infinitely many # and arbitrarily long input blocks, then $\binom{\alpha(0)}{\beta(0)}\binom{\alpha(1)}{\beta(1)}\binom{\alpha(2)}{\beta(2)}\cdots$ contains arbitrarily long output blocks. It is easy to come up with a WMSO+U formula defining L by formalizing the definitions of input and output blocks in first-order logic.

Now, consider L as winning condition for a delay game. Intuitively, Player O has to specify arbitrarily long output blocks, provided Player I produces arbitrarily long input blocks. The challenge for Player O is that she has to specify at the beginning of every output block whether she ends the block in a position where Player O has picked a 0 or a 1.

First, we show that Player O wins $\Gamma_f(L)$ for the delay function f with $f(i) = 2$ for every i. Consider round i for some i with $\alpha(i) = \#$. By the choice of f, Player O has already picked $\alpha(i+1)\cdots\alpha(2i+1)$. Let j in the range $i+1 \leq j \leq 2i+1$ be maximal with $\alpha(i+1)\cdots\alpha(j) \in \{0,1\}^+$. If j is defined, then Player O picks $\alpha(j)$ in round i, $*$ during rounds $i+1, \ldots, j-1$, and $\alpha(j)$ in round j. In every other situation, she picks an arbitrary letter.

Now, consider an outcome $\binom{\alpha(0)}{\beta(0)}\binom{\alpha(1)}{\beta(1)}\binom{\alpha(2)}{\beta(2)}\cdots$ that is consistent with this strategy that contains infinitely many # and arbitrarily long input blocks. Let $n \in \mathbb{N}$ be arbitrary and pick an input block of length at least n. We can assume w.l.o.g. that the block begins at $\alpha(i)$ with $i > n$. Thus, in round i when Player O had to pick $\beta(i)$, Player I had already picked $\alpha(i+1)\cdots\alpha(2i+1)$, which comprises the whole input block of length n. Accordingly, Player O produces an output block of length at least n. Thus, the winning condition is satisfied.

It remains to show that Player I wins $\Gamma_f(L)$ for every constant delay function f. He uses a counter c to produce arbitrarily long input blocks, which is initialized to $f(0)$. In round 0, he picks $\#0^{f(0)-1}$. If Player O answers the # with 0, then Player I continues picking 1 until he has produced an input block of length c. Dually, if Player O answers the # with 1 or $*$, then Player I continues picking 0 until he has produced an input block of length c. In both cases, he continues by incrementing the counter and by picking $\#0^{f(0)-1}$ during the next rounds until Player O has to respond to the last #. Then, he continues as in the case distinction described above.

Now, consider an outcome $\binom{\alpha(0)}{\beta(0)}\binom{\alpha(1)}{\beta(1)}\binom{\alpha(2)}{\beta(2)}\cdots$ that is consistent with this strategy. It contains infinitely many # and arbitrarily long input blocks. Fur-

thermore, the length of every output block is bounded by $f(0)$, as Player I is able to react to Player O's declaration at the beginning of each such block by playing the opposite letter. Thus, the play is winning for Player I. □

We have just shown that Player O wins the game for the delay function $f(i) = 2$. Such a function is called linear [15], as the lookahead grows linearly.

6 Conclusion

We considered delay games with max-regular winning conditions. Our main result is an algorithm that determines whether Player O has a winning strategy for some constant delay function, which consists of reducing the original problem to a delay-free game with max-regular winning condition. Such a game can be solved by encoding it as an emptiness problem for a certain class of tree automata (so-called WMSO+UP automata) that capture WMSO+UP on infinite trees. Our reduction also yields a doubly-exponential upper bound on the necessary constant delay to win such a game, provided Player O does win for some constant delay function.

It is open whether the doubly-exponential upper bound is tight. The best lower bounds are exponential and hold already for deterministic reachability and safety automata [20], which can easily be transformed into max-automata.

We deliberately skipped the complexity analysis of our algorithm, since the reduction of the delay-free game to an emptiness problem for WMSO+UP automata does most likely not yield tight upper bounds on the complexity. Instead, we propose to investigate (delay-free) games with max-regular winning conditions, a problem that is worthwhile studying on its own, and to find a direct solution algorithm. Currently, the best lower bound on the computational complexity of determining whether Player O wins a delay game with max-regular winning condition for some constant delay function is the EXPTIME-hardness result for games with safety conditions [20].

Also, we showed that constant delay is not sufficient for max-regular conditions by giving a condition L such that Player O wins $\Gamma_f(L)$ for some linear f, but not for any constant delay function f.

Both the lower bound on the necessary lookahead and the one on the computational complexity for safety conditions mentioned above are complemented by matching upper bounds for games with parity conditions [20], i.e., having a parity condition instead of a safety condition has no discernible influence. Stated differently, the complexity of the problems manifests itself in the transition structure of the automaton. Our example from Section 5 shows that this is no longer true for max-regular conditions: having a quantitative acceptance condition requires growing lookahead.

In ongoing work, we aim to solve delay game with respect to arbitrary delay functions and to determine whether linear delay functions are sufficient to win delay games with max-regular winning conditions.

References

1. Bojańczyk, M.: Weak MSO with the unbounding quantifier. Theory Comput. Syst. 48(3), 554–576 (2011)
2. Bojańczyk, M.: Weak MSO+U with path quantifiers over infinite trees. In: Esparza et al. [12], pp. 38–49
3. Bojańczyk, M., Gogacz, T., Michalewski, H., Skrzypczak, M.: On the decidability of MSO+U on infinite trees. In: Esparza et al. [12], pp. 50–61
4. Bojańczyk, M., Parys, P., Toruńczyk, S.: The MSO+U theory of $(\mathbb{N}, <)$ is undecidable (2015), arXiv 1502.04578
5. Bojańczyk, M., Toruńczyk, S.: Weak MSO+U over infinite trees. In: Dürr, C., Wilke, T. (eds.) STACS 2012. LIPIcs, vol. 14, pp. 648–660. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2012)
6. Büchi, J.R.: On a Decision Method in Restricted Second-Order Arithmetic. In: International Congress on Logic, Methodology, and Philosophy of Science. pp. 1–11. Stanford University Press (1962)
7. Büchi, J.R., Landweber, L.H.: Solving sequential conditions by finite-state strategies. Trans. Amer. Math. Soc. 138, pp. 295–311 (1969)
8. Carayol, A., Löding, C.: MSO on the infinite binary tree: Choice and order. In: Duparc, J., Henzinger, T.A. (eds.) CSL 2007. LNCS, vol. 4646, pp. 161–176. Springer (2007)
9. Carayol, A., Löding, C.: Uniformization in automata theory. In: Schroeder-Heister, P., Heinzmann, G., Hodges, W., Bour, P.E. (eds.) International Congress of Logic, Methodology and Philosophy of Science. College Publications, London (2012), to appear
10. Chatterjee, K., Henzinger, T.A., Horn, F.: Finitary winning in omega-regular games. ACM Trans. Comput. Log. 11(1) (2009)
11. Emerson, E.A., Jutla, C.S.: Tree automata, mu-calculus and determinacy (extended abstract). In: FOCS 1991. pp. 368–377. IEEE (1991)
12. Esparza, J., Fraigniaud, P., Husfeldt, T., Koutsoupias, E. (eds.): ICALP 2014, Part II, LNCS, vol. 8573. Springer (2014)
13. Faymonville, P., Zimmermann, M.: Parametric linear dynamic logic. In: Peron and Piazza [26], pp. 60–73
14. Fijalkow, N., Zimmermann, M.: Parity and Streett Games with Costs. LMCS 10(2) (2014)
15. Fridman, W., Löding, C., Zimmermann, M.: Degrees of Lookahead in Context-free Infinite Games. In: Bezem, M. (ed.) CSL 2011. LIPIcs, vol. 12, pp. 264–276. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2011)
16. Grädel, E., Thomas, W., Wilke, T. (eds.): Automata, Logics, and Infinite Games: A Guide to Current Research, LNCS, vol. 2500. Springer (2002)
17. Gurevich, Y., Shelah, S.: Rabin's uniformization problem. The Journal of Symbolic Logic 48, 1105–1119 (1983)
18. Holtmann, M., Kaiser, L., Thomas, W.: Degrees of lookahead in regular infinite games. LMCS 8(3) (2012)
19. Hosch, F.A., Landweber, L.H.: Finite delay solutions for sequential conditions. In: ICALP 1972. pp. 45–60 (1972)
20. Klein, F., Zimmermann, M.: How much lookahead is needed to win infinite games? In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) ICALP 2015 (Part 2). LNCS, vol. 9135, pp. 452–463. Springer (2015)

21. Klein, F., Zimmermann, M.: What are Strategies in Delay Games? Borel Determinacy for Games with Lookahead. In: Kreutzer, S. (ed.) CSL 2015. LIPIcs, vol. 41, pp. 519–533. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2015)
22. Kupferman, O., Piterman, N., Vardi, M.Y.: From liveness to promptness. Formal Methods in System Design 34(2), 83–103 (2009)
23. Löding, C., Winter, S.: Synthesis of deterministic top-down tree transducers from automatic tree relations. In: Peron and Piazza [26], pp. 88–101
24. Martin, D.A.: Borel determinacy. Annals of Mathematics 102, 363–371 (1975)
25. Mostowski, A.: Games with forbidden positions. Tech. Rep. 78, University of Gdańsk (1991)
26. Peron, A., Piazza, C. (eds.): GandALF 2014, EPTCS, vol. 161 (2014)
27. Thomas, W.: Infinite games and uniformization. In: Banerjee, M., Seth, A. (eds.) ICLA. LNCS, vol. 6521, pp. 19–21. Springer (2011)
28. Thomas, W., Lescow, H.: Logical specifications of infinite computations. In: de Bakker, J.W., de Roever, W.P., Rozenberg, G. (eds.) A Decade of Concurrency, Reflections and Perspectives, REX School/Symposium. LNCS, vol. 803, pp. 583–621. Springer (1993)
29. Trakhtenbrot, B., Barzdin, I.: Finite Automata; Behavior and Synthesis. Fundamental Studies in Computer Science, V. 1, North-Holland Publishing Company; New York: American Elsevier (1973)
30. Zimmermann, M.: Optimal Bounds in Parametric LTL Games. Theoret. Comput. Sci. 493(0), 30 – 45 (2013)
31. Zimmermann, M.: Delay games with WMSO+U winning conditions. In: Beklemishev, L.D., Musatov, D.V. (eds.) CSR 2015, LNCS, vol. 9139, pp. 412–425. Springer (2015)